Hardware-Based Linear Program Decoding with the Alternating Direction Method of Multipliers

Mitchell Wasson, Mario Milicevic, *Student Member, IEEE*, Stark C. Draper, *Senior Member, IEEE*, and Glenn Gulak, *Senior Member, IEEE*

arXiv:1611.05975v1 [cs.IT] 18 Nov 2016

***Abstract*—We present a hardware-based implementation of Linear Program (LP) decoding for binary linear codes. LP decoding frames error-correction as an optimization problem. In contrast, variants of Belief Propagation (BP) decoding frame error-correction as a problem of graphical inference. LP decoding has several advantages over BP-based methods, including convergence guarantees and better error-rate performance in high-reliability channels. The latter makes LP decoding attractive for optical transport and storage applications. However, LP decoding, when implemented with general solvers, does not scale to large blocklengths and is not suitable for a parallelized implementation in hardware. It has been recently shown that the Alternating Direction Method of Multipliers (ADMM) can be applied to decompose the LP decoding problem. The result is a message-passing algorithm with a structure very similar to BP. We present new intuition for this decoding algorithm as well as for its major computational primitive: projection onto the parity polytope. Furthermore, we present results for a fixed-point Verilog implementation of ADMM-LP decoding. This implementation targets a Field-Programmable Gate Array (FPGA) platform to evaluate error-rate performance and estimate resource usage. We show that Frame Error Rate (FER) performance well within 0.5dB of double-precision implementations is possible with 10-bit messages. Finally, we outline a number of research opportunities that should be explored en-route to the realization of an Application Specific Integrated Circuit (ASIC) implementation capable of gigabit per second throughput.**

I. Introduction

THE field of error-correction coding was revolutionized in the mid-1990s by the widespread adoption (and academic study) of graph-based codes and associated Belief Propagation (BP) message-passing decoding algorithms [1]–[3]. A key aspect of the success of these codes was their compatibility with hardware. BP-based decoders are naturally distributed algorithms and variants such as Min-Sum are (relatively) easily mapped to hardware. Graph-based codes, particularly Turbo codes and Low-Density Parity-Check (LDPC) codes, have been adopted in many real world systems. However, there are issues present with BP-based decoding algorithms. The first is their reliance on the tree assumption for the code-defining graph. In practice tree codes are not used due to their poor distance properties. This results in the use of LDPC codes without performance or convergence guarantees due to graph cycles. Additionally, it is observed in practice that BP-based decoding algorithms often suffer from performance deficiencies, termed "error floors", in high-reliability channels [4].

In the early 2000s, Feldman and his collaborators realized that the Maximum Likelihood (ML) decoding problem for binary linear codes can be rephrased as an integer program [5]. One obtains a Linear Program (LP) by relaxing the integer constraints. Feldman's work applies to any binary linear code, but he concentrated on LDPC codes due to their prevalence and smaller constraint sets. These results generated much interest among coding theorists. LPs are an extremely well-studied and understood class of optimization problems, especially when contrasted with BP. For instance, LP decoding has an ML certificate property [5], such that if LP decoding fails, it fails in a detectable way (to a non-integer vertex). The relaxation can then be tightened and the LP re-run [6]. If a high-quality expander or high-girth code is used, LP decoding is guaranteed to correct a constant number of bit flips [7], [8]. Broadly, it was hoped that by studying LP decoding, more would be understood about BP decoding.

On the practical side, there was less excitement. There initially seemed to be no real-world need for such a decoder. Furthermore, traditional LP solvers did not scale easily to the blocklengths of modern error-correcting codes. Nevertheless, a number of groups did study how to build an application-specific low-complexity LP decoder [6], [9]–[11]. In particular, Barman et al. [11] built an application-specific LP decoder that was computationally competitive with BP and that had a message-passing structure with a standard schedule [11]. They solved the LP decoding problem using the Alternating Direction Method of Multipliers (ADMM), a decomposition technique used in large-scale optimization [12]. Able to study LP decoding performance at long blocklengths, it was observed empirically, and later confirmed theoretically, that LP decoding far outperforms BP in the high Signal-to-Noise Ratio (SNR) regime [11], [13], [14]. In this regime, LP decoders do not suffer from the same error floor effects as BP. Using the ADMM solver, Liu and Draper were able to augment the objective of LP decoding with a penalty term to improve error-rates further in low-reliability channels [15]. Additionally, LP decoding can be used as a subroutine in a multi-stage decoder that quickly approaches ML performance [16]. Thus, for applications in which reliability demands are extreme, LP decoding is an attractive alternative (or complement) to BP. Further generalization of ADMM-LP to non-binary and

This material was presented in part at the 2015 Asilomar Conf. on Signals, Systems, and Computers, Pacific Grove, CA, Nov. 2015.

M. Wasson was with, and M. Milicevic, S. Draper, and G. Gulak are with, the Dept. of Electrical and Comp. Eng., University of Toronto, ON M5S 3G4, Canada (e-mail: m.wasson@mail.utoronto.ca, mario.milicevic@utoronto.ca, stark.draper@utoronto.ca, gulak@eecg.toronto.edu)

This work was supported by the National Science Foundation (NSF) under Grant CCF-1217058, by the Natural Science and Engineering Research Council (NSERC) of Canada, including through a Discovery Research Grant, and by the Canadian Microelectronics Corporation (CMC).

multipermutation codes are developed in [17], [18].

In parallel to these theoretical and algorithmic developments, there has been growing interest in moving ADMM-LP toward a hardware implementation. Several groups have made progress in creating efficient methods for solving the key computational primitive of ADMM-LP decoding, Euclidean projection onto the "parity polytope" [19]–[21]. In particular, Wasson and Draper investigated mapping this operation to hardware [21]. Several implementation papers have also considered ADMM-LP decoding in other contexts. Debbabi et al. investigated how to schedule messages more efficiently and developed a multicore implementation [22], [23]. Jiao et al. modified penalization to improve error-rate performance of irregular LDPC codes [24]. Finally, Wei et al. implemented ADMM-LP avoiding projections when possible [25].

While useful investigations, these studies do not demonstrate whether or not ADMM-LP decoding is viable in hardware. In this paper, we present a Field-Programmable Gate Array (FPGA)-based implementation that shows that the ADMM-LP decoding algorithm can be mapped to hardware without suffering an unacceptable performance loss. First, we review and expand upon the execution and intuition of the ADMM-LP decoding algorithm. Next, we review the developments made in [21] to implement Euclidean projection onto the parity polytope in hardware. We then describe how to assemble the pieces to form a complete LP decoder. We present results for the [155, 64] Quasi-Cyclic (QC) LDPC code introduced by Tanner et al. [26], the [672, 546] QC-LDPC code from the IEEE 802.11ad (WiGig) standard [27], and an ensemble of (3,6)-regular [1002, 503] QC-LDPC codes. We test code performance using an FPGA-based simulation environment. While our initial implementation requires more hardware resources than Min-Sum decoders, we find that it is possible to preserve the superior error-rate performance of ADMM-LP in fixed point.

II. Background

In this paper, we consider the decoding of binary linear codes. A binary linear code $\mathcal{C}$ of blocklength n is a k-dimensional subspace of $\mathbb{F}_2^n$. Such a code can be defined as the null space of the $m \times n$ "parity-check" matrix H, i.e., $\mathcal{C} = \{x \in \{0,1\}^n : Hx = 0 \pmod 2\}$. In general $m \geq n-k$ with equality when H has full rank. The rate of $\mathcal{C}$ is defined to be $R = k/n$ which specifies the number of information bits transmitted per codeword symbol. Each row of the parity-check matrix corresponds to a check, which specifies a subset of codeword symbols that must add to 0 modulo 2. These checks are indexed by the set $\mathcal{J} = \{1, \ldots, m\}$. Each column of the parity-check matrix corresponds to a codeword symbol or variable, indexed by $\mathcal{I} = \{1, \ldots, n\}$. The neighborhood of check j, denoted $\mathcal{N}_c(j)$, is the set of variables that check j constrains to add to 0. That is, $\mathcal{N}_c(j) = \{i \in \mathcal{I} : H_{j,i} = 1\}$. Similarly, the neighborhood of variable i, $\mathcal{N}_v(i)$, is the set of checks in which variable i participates, $\mathcal{N}_v(i) = \{j \in \mathcal{J} : H_{j,i} = 1\}$.

Given a stochastic channel model $\mathcal{P}(y|x)$ where $y \in \mathcal{Y}^n$ is the channel output, ML decoding amounts to maximizing the model over the set of codewords. That is, we decode to $\arg\max_{x \in \mathcal{C}} \mathcal{P}(y|x)$. It was shown in [28] that, when considering a binary linear code transmitted over a symmetric memoryless channel, the ML decoding objective is linear in the length-n vector γ of Log-likelihood Ratios (LLRs) $\gamma_i = \log(\mathcal{P}(y_i|0)/\mathcal{P}(y_i|1))$. ML decoding problem thus is

$$\arg\max_{x \in \mathcal{C}} \mathcal{P}(y|x) = \arg\min_{x \in \mathcal{C}} \sum_{i=1}^{n} \gamma_i x_i = \arg\min_{x \in \mathcal{C}} \gamma^T x. \quad (1)$$

We note that γ can be multiplied by any positive scalar without changing the problem.

Having framed ML as an optimization problem with a linear objective, we are ready to develop the LP relaxation first proposed in [28]. First, denote by $x_{\mathcal{S}}$, $\mathcal{S} \subseteq \mathcal{I}$, the length-$|\mathcal{S}|$ vector formed with the components of x indexed by $\mathcal{S}$. With this notation, we can restate the parity-check condition for a valid codeword as $\mathcal{C} = \left\{x \in \{0,1\}^n : 1^\top x_{\mathcal{N}_c(j)} = 0 \pmod 2 \text{ for all } j \in \mathcal{J}\right\}$. Each of the m constraints in this set can be visualized as requiring that the set of codeword variables connected to any particular check must be an even-weight vertex of the unit hyper-cube.

The LP decoding problem results from relaxing these constraints [5], [28]. Instead of requiring the vector of variables connected to each check to be an even-weight vertex of the unit hyper-cube, LP decoding rather requires this set of variables to lie in the convex hull of the vertices. Visualized in Fig. 1, the convex hull of the even-weight vertices of the unit hyper-cube is termed the "parity polytope", denoted $\mathbb{PP}_d$ in d-dimension:

$$\mathbb{PP}_d := \text{conv}\left(\left\{e \in \{0,1\}^d : 1^\top e = 0 \pmod 2\right\}\right). \quad (2)$$

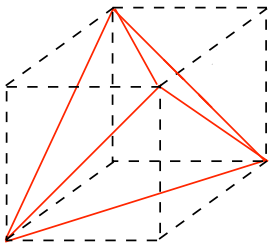

Fig. 1. Visualization of $\mathbb{PP}_3$ within the three-dimensional unit cube.

The polytope $\mathbb{PP}_d$ can be explicitly defined by a number of half-space inequalities [28]. Every odd-weight vertex is surrounded by even-weight vertices. Each half-space inequality is defined by the hyperplane that contains all these even-weight vertices and "cuts" off the half of the space in which the odd-weight vertex sits. Half of the 2^d vertices are of odd weight, so we can describe $\mathbb{PP}_d$ with 2^{d-1} half-space constraints. Each such inequality corresponds to one of the constraints in the first line of the following description of $\mathbb{PP}_d$ where we use the notation $[d] = \{1, \ldots, d\}$. A vector $v \in \mathbb{PP}_d$ if

$$\begin{aligned} &\sum_{i \in S} v_i - \sum_{i \in [d] \setminus \mathcal{S}} v_i \leq |\mathcal{S}| - 1 && \mathcal{S} \subseteq [d],\ |\mathcal{S}| \text{ odd} \\ &0 \leq v_i \leq 1 && i \in [d] \end{aligned} \quad . \quad (3)$$

The box constraints $0 \leq v_i \leq 1$ are not always redundant, e.g., when $d = 2$.

In summary, LP decoding requires us to solve

$$\begin{array}{rl} \arg\min\limits_{x} & \gamma^\top x \\ \text{subject to} & x_{\mathcal{N}_c(j)} \in \mathbb{PP}_{|\mathcal{N}_c(j)|} \quad j \in \mathcal{J} \\ & x \in [0,1]^n \end{array} \tag{4}$$

Note that LP decoding is not guaranteed to yield the ML solution. Due to the relaxation, the feasible space has fractional vertices. The failure model of LP decoding is when one of these "pseudocodewords" is the minimal cost vertex.

One might think of rounding the fractional components when the LP solver outputs a pseudocodeword. However, this does not solve the pseudocodewords problem. An alternative approach proposed by Liu and Draper [15] is to augment the objective of (4) with a penalty function. This approach, referred to as "penalized LP" decoding, discourages pseudocodewords by penalizing the closeness of variables to $1/2$. Many penalty functions were tested in [15], but we only implement the so-called ℓ_1-penalty function, due to its good error-rate performance and algorithmic simplicity. The ℓ_1-penalized LP decoding problem is given by

$$\begin{array}{rl} \min & \gamma^\top x - \alpha\left\|x - \frac{1}{2}\right\|_1 \\ \text{subject to} & x_{\mathcal{N}_c(j)} \in \mathbb{PP}_{|\mathcal{N}_c(j)|} \quad j \in \mathcal{J} \\ & x \in [0,1]^n \end{array} \tag{5}$$

where $\alpha \geq 0$ is termed the penalty parameter. The penalty parameter tunes how severely non-binary variables should be penalized. Setting $\alpha = 0$ reduces (5) to (4). While moderate values of α improve performance in the waterfall regime, an excessively large α can adversely affect performance [15].

Up until this point, we have been discussing prior formulations of LP decoding. We now discuss a simple transformation of LP decoding that proves useful when designing a fixed-point implementation. The original LP decoding formulation operates inside the unit hypercube centered around $1/2$. In our hardware implementation, signed integers are used to implement fixed-point arithmetic. Therefore, to eliminate possible asymmetries, we prefer LP decoding to operate inside the unit hypercube symmetrically centered around 0. To accomplish this, the simple variable substitution $x_{\text{new}} = x_{\text{old}} - 1/2$ can be applied to (5). The result

$$\begin{array}{rl} \min & \gamma^\top x - \alpha\|x\|_1 \\ \text{subject to} & x_{\mathcal{N}_c(j)} \in \mathbb{PP}_{|\mathcal{N}_c(j)|} - \frac{1}{2} \quad j \in \mathcal{J} \\ & x \in [-1/2, 1/2]^n \end{array} \tag{6}$$

is an equivalent optimization problem with two important differences. The first is that the objective now penalizes closeness to 0 rather than to $1/2$. The second is that check neighborhoods must be in the "centered" parity polytope. The d-dimensional centered parity polytope $\mathbb{PP}_{|\mathcal{N}_c(j)|} - \frac{1}{2}$ is obtained by taking every point in $\mathbb{PP}_d$ and subtracting the length-d all $1/2$ vector. For simplicity we subsequently refer to this shifted object simply as the parity polytope, unless disambiguation is required.

III. Algorithms

In Section III-A, we discuss the ADMM algorithm, its application to error-correction decoding and message passing interpretation. The two key subroutines: projection onto the parity polytope and projection onto the probability simplex are discussed in Sections III-B and III-C, respectively.

A. ADMM Decomposition and Message Passing

The characteristic that, in a linear code, each component of the codeword estimate x (generally) participates in multiple check constraints inhibits the decomposability of LP decoding. A small modification is therefore introduced in [11] to apply ADMM to LP decoding . We define an auxiliary "replica" variable vector $z_j = x_{\mathcal{N}_c(j)}$ for each check neighborhood. By substituting into (6), we arrive at the following result, which fits the ADMM template:

$$\begin{array}{rll} \min & \gamma^\top x - \alpha\|x\|_1 & \\ \text{subject to} & x \in [-\frac{1}{2}, \frac{1}{2}]^n & \\ & z_j \in \mathbb{PP}_{|\mathcal{N}_c(j)|} - \frac{1}{2} & j \in \mathcal{J} \\ & z_j = x_{\mathcal{N}_c(j)} & j \in \mathcal{J}. \end{array} \tag{7}$$

The ADMM decomposition for (penalized) LP decoding starts from the ℓ_2-regularized Lagrangian

$$\begin{aligned} L_\mu(x,z,\lambda) = \; & \gamma^\top x - \alpha\|x\|_1 + \sum_{j\in\mathcal{J}} \lambda_j^\top \left(x_{\mathcal{N}_c(j)} - z_j\right) \\ & + \tfrac{\mu}{2} \sum_{j\in\mathcal{J}} \left\|x_{\mathcal{N}_c(j)} - z_j\right\|_2^2 \end{aligned}.$$

We use z and λ to refer to the z_j and λ_j in aggregate. The λ_j are length-$|\mathcal{N}_c(j)|$ dual variable vectors that enforce the $z_j = x_{\mathcal{N}_c(j)}$ equality constraints. The μ parameter is a positive number that determines the degree of regularization. While regularization does not change the solution of the optimization problem, it accelerates algorithmic convergence [12].

ADMM-LP decoding alternates, in a round-robin manner, between minimizing over codeword estimates x and replicas z, followed by an update of the dual variables λ. Letting $\mathcal{X}$ and $\mathcal{Z}$ represent the feasible sets of x and z (the dual variables are unconstrained), each iteration takes the form [11], [15]

$$\begin{array}{rll} x \leftarrow & \arg\min_{x\in\mathcal{X}} L_\mu(x,z,\lambda) & \\ z \leftarrow & \arg\min_{z\in\mathcal{Z}} L_\mu(x,z,\lambda) & \\ \lambda_j \leftarrow & \lambda_j + \mu\left(x_{\mathcal{N}_c(j)} - z_j\right) & j \in \mathcal{J}. \end{array} \tag{8}$$

The x update can be decomposed into individual variable updates since the solution to its optimization problem separates into distinct calculations for each variable [11], [15]. Similarly, the z update can be decomposed to update each z_j individually. The λ update is already expressed in a decomposed from. When these update rules are fleshed out, μ can be eliminated by reparameterizing γ, α and λ by a factor of μ [29]. Therefore μ is not included in the ADMM-LP decoding algorithm statement and the values of γ, α and λ have slightly different parameterizations than in (8) going forward.

The fact that the updates decompose means that the algorithm performs a set of parallel variable updates followed by a set of parallel check updates. The result is a message-passing algorithm with a structure similar to BP. Variable update i is performed using the Log-Likelihood Ratio (LLR) γ_i and a length-$|\mathcal{N}_v(i)|$ vector of messages from each of its neighboring checks, denoted $m_{\mathcal{N}_v(i)\to i}$. Check update j is performed using the dual variable vector λ_j and a length-$|\mathcal{N}_c(j)|$ vector of

Alg. 1. Given a LLR vector $\gamma \in \mathbb{R}^n$, calculate the ℓ_1-penalized LP decoding from (7) using ADMM. The current estimate x is returned upon termination.

1: Initialize all λ_j's and $m_{\mathcal{N}_v(i)\to i}$'s to 0.
2: **repeat**
3: **for** $i \in \mathcal{I}$ **do** ▷ Variable Updates
4: $t_i \leftarrow 1^\top m_{\mathcal{N}_v(i)\to i} - \gamma_i$
5: $s_i \leftarrow \begin{cases} t_i + \alpha & \text{if } t_i > 0 \\ t_i & \text{if } t_i = 0 \\ t_i - \alpha & \text{if } t_i < 0 \end{cases}$
6: $x_i \leftarrow \prod_{[-\frac{1}{2},\frac{1}{2}]} \left(\frac{s_i}{|\mathcal{N}_v(i)|} \right)$
7: **end for**
8: **for** $j \in \mathcal{J}$ **do** ▷ Check Updates
9: $v_j \leftarrow x_{\mathcal{N}_c(j)} + \lambda_j$
10: $z_j \leftarrow \prod_{\mathbb{PP}_{|\mathcal{N}_c(j)|}-\frac{1}{2}} (v_j)$
11: $\lambda_j \leftarrow v_j - z_j$
12: $m_{j\to\mathcal{N}_c(j)} \leftarrow 2z_j - v_j$
13: **end for**
14: **until** termination

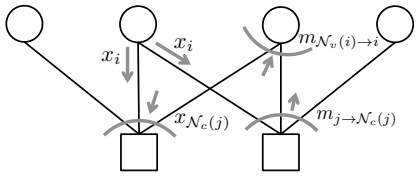

$$H = \begin{pmatrix} 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 \end{pmatrix}$$

Fig. 2. A parity-check matrix and associated factor graph with messages labelled. Checks vertices are drawn as squares.

messages from each of the neighboring variables. The latter contains the current estimates $x_{\mathcal{N}_c(j)}$ of neighboring variables. The result of the update of check j is a new estimate of the associated dual variables λ_j as well as a length-$|\mathcal{N}_v(i)|$ message vector whose components are sent to neighboring variables. This vector is denoted $m_{j\to\mathcal{N}_c(j)}$. Alg. 1 presents the ADMM-LP decoding algorithm in full. The notation $\prod_{\mathcal{A}}(\cdot)$ denotes Euclidean projection onto the set $\mathcal{A}$.

Viewing the variable/constraint structure of ADMM-LP in (8) using a factor graph, we observe that the message-passing schedule followed in Alg. 1 is the standard flooding schedule of BP. This view helps to highlight key differences between ADMM-LP and BP. A factor graph is a bipartite graph whose connectivity, when representing a linear code, is specified by the code's parity-check matrix [3]. For each codeword variable there is a variable vertex. For each parity check there is a factor (or check) vertex. Variable vertex i and check vertex j are connected if $H_{j,i} = 1$. The execution of ADMM-LP can be viewed as variable updates taking place inside variable vertices and check updates taking place inside check vertices. Messages are passed between variables and checks along the edges of the graph. To illustrate this interpretation, Fig. 2 depicts a parity-check matrix and associated factor graph; edges are labeled by ADMM-LP messages.

In Fig. 2 the $m_{\mathcal{N}_v(i)\to i}$ messages are composed of vector components from the $m_{j\to\mathcal{N}_c(j)}$ messages. The λ_j do not appear in the figure. The vector of dual variables λ_j is used only within the j^{th} check update. Dual variable vectors serve as internal check states that are not passed as messages. This is an important difference from BP as the dual variables play an important role in improving error-floor performance [14]. A second difference is that the $|\mathcal{N}_v(i)|$ outgoing messages from any variable i are all equal, corresponding to the current estimate of x_i. This differs from BP where, in general, all outgoing messages from a variable will differ.

We now examine the steps of Alg. 1. On line 4, variable updates first sum incoming messages and the negative LLR to form a variable estimate. Incoming messages tell the variable what value it should take on. Next, on line 5, the penalization is applied. A non-zero penalty pushes each variable estimate in the direction of its current belief. Recall that this is done to discourage fractional solutions (pseudocodewords). When α is small, the effect of penalization is reduced, making the algorithm closer to (unpenalized) LP decoding. A slight difference in Alg. 1 from penalized LP decoding's original derivation in [15] is that no penalty is applied if $t_i = 0$. This modification is important in a fixed-point implementation to avoid bias in codeword estimates. On line 6, the penalized estimate is normalized by the variable degree and projected onto the $[-\frac{1}{2}, \frac{1}{2}]$ interval. The resulting final estimate is then passed to neighboring checks. Roughly speaking, the variable estimate is the average of the incoming messages. On line 9, the first step in the check update is to take the vector of neighboring variable estimates and add the vector of dual variables (the check state vector). An updated vector of the replica estimate is obtained by projecting the addition result v_j onto the parity polytope. This is where the parity polytope constraints of LP decoding are enforced. Using the projection, a new check state λ_j and set of outgoing messages $m_{j\to\mathcal{N}_c(j)}$ are calculated (possibly in parallel) on lines 11 and 12.

We think of the dual variable estimates as affecting algorithmic progression in two major ways. First, λ_j acts as a momentum term on line 9. It brings v_j closer to the previous value of v_j, ensuring that z_j does not evolve too erratically. Second, according to the pricing interpretation of duality, the λ_j specifies the cost of breaking the equality constraint $z_j = x_{\mathcal{N}_c(j)}$. We can see this effect more clearly if line 12 is rewritten as $m_{j\to\mathcal{N}_c(j)} \leftarrow z_j - (v_j - z_j) = z_j - \lambda_j$. Since the new λ_j value is the mismatch between v_j and z_j, line 12 compensates this mismatch by including it in the outgoing messages. At convergence, the λ_j subtracted off here is canceled by the λ_j added in to compute v_j on line 9.

Note that a termination condition is not specified in Alg. 1. While algorithmic convergence can be used as the stopping criterion in floating point [11], it may not be possible to obtain convergence in fixed point to an arbitrary precision. Thus, in our implementation, we impose a fixed number of iterations, but can also terminate early if rounding the current codeword estimate produces a codeword.

While their message-passing schedules are the same, we have already observed two differences between ADMM-LP and BP: the existence of dual variables that form the check states, and the fact that all outgoing messages from a variable node are identical. To this point, a third significant difference

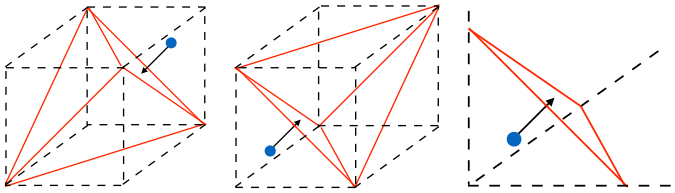

(a) Identify active facet (b) Transform problem (c) Simplex projection

Fig. 3. Projection onto the parity polytope $\mathbb{PP}_3 - \frac{1}{2}$: identify active facet, transform into canonical coordinate system, project onto probability simplex.

has been abstracted: the computational primitive of the check update, which is the Euclidean projection onto the parity polytope. A discussion of how to implement this projection efficiently in hardware is the topic of the next section.

B. Parity Polytope Projection

Euclidean projection of a vector v onto the d-dimensional parity polytope is specified by the quadratic program

$$\prod_{\mathbb{PP}_d - \frac{1}{2}}(v) = \arg \min_{w \in \mathbb{PP}_d - \frac{1}{2}} \|w - v\|_2^2. \tag{9}$$

Projection onto the centered $(\mathbb{PP}_d - \frac{1}{2})$ and non-centered $(\mathbb{PP}_d)$ parity polytope are similar operations related as

$$\prod_{\mathbb{PP}_d - \frac{1}{2}}(v) = \prod_{\mathbb{PP}_d}\left(v + \frac{1}{2}\right) - \frac{1}{2}.$$

Barman et al. began the investigation into efficient projection onto the parity polytope [11], [30]. These researchers established a "two-slice" representation of the polytope and exploited rotational symmetry to sort the components of v into a canonical coordinate system for projection (and subsequent de-sort). However, their algorithm is not well-suited for hardware due to its iterative nature and complexity of the sorting procedure. X. Zhang and Siegel [19], [31] improved the method by removing the sort and de-sort operations through efficient identification of the violated cut from (3). Unfortunately, as with the first approaches, the method remains intensively iterative. In parallel to [19], G. Zhang et al. made the connection to projection onto the probability simplex [20] which provides clean geometric intuition. In [21] Wasson and Draper combined the advances of [19], [20] to create the hardware-compatible method of projection we now describe.

First, the violated cut from (3) is identified, revealing the active facet. The identified cut defines a similarity transform used to reorient the problem into a canonical orientation. The problem is thereby reduced to projection onto the (centered) probability simplex. After projecting onto the simplex, the similarity transform is inverted to yield the projection onto the parity polytope. This high-level description is depicted in Fig. 3. The algorithm described in Alg. 2 was slightly modified in [29] from the algorithm presented in [21] to project onto the centered parity polytope. The algorithm has a straightforward, non-iterative, execution path whose steps can largely be parallelized. This, combined with simple intuition, makes Alg. 2 an excellent candidate for hardware adoption.

Alg. 2. Given a vector $v \in \mathbb{R}^d$, calculate its Euclidean projection onto the parity polytope. $w = \prod_{\mathbb{PP}_d - \frac{1}{2}}(v)$ is returned from the method.

1: **for** $i \in [d]$ **do** ▷ Facet Identification
2: $\quad f_i \leftarrow \begin{cases} 1 & \text{if } v_i \geq 0 \\ 0 & \text{otherwise} \end{cases}$
3: **end for**
4: **if** $1^\top f$ is even **then**
5: $\quad i^* \leftarrow \arg\min_{i \in [d]} |v_i|$
6: $\quad f_{i^*} \leftarrow 1 - f_{i^*}$
7: **end if**
8: **for** $i \in [d]$ **do** ▷ Similarity Transform
9: $\quad \tilde{v}_i \leftarrow v_i(-1)^{f_i}$
10: **end for**
11: $\tilde{u} \leftarrow \prod_{\mathbb{S}_d - \frac{1}{2}}(\tilde{v})$ ▷ Simplex Projection
12: **for** $i \in [d]$ **do** ▷ Similarity Transform
13: $\quad u_i \leftarrow \tilde{u}_i(-1)^{f_i}$
14: **end for**
15: **if** $1^\top \prod_{[-\frac{1}{2}, \frac{1}{2}]^d}(\tilde{v}) \geq 1 - \frac{d}{2}$ **then** ▷ Membership Test
16: $\quad w \leftarrow \prod_{[-\frac{1}{2}, \frac{1}{2}]^d}(v)$
17: **else**
18: $\quad w \leftarrow u$
19: **end if**

In Alg. 2, lines 1–7 form the facet identification portion of the projection algorithm. The objective here is to identify the vertex cut from (3) that is violated (if one is violated). This amounts to finding the closest odd-weight vertex of the unit hypercube [32]. First, the closest vertex of the hypercube is found and stored in the binary vector f. For the non-centered case considered in Alg. 2, the actual vertex is $f - 0.5$. If the Hamming weight of f is odd, then the closest vertex violates the parity constraint (i.e., it is *not* a codeword of the single parity-check code) and we have identified the violated cut. On the other hand, if the Hamming weight of f is even, the nearest vertex does not violate the parity constraint. The Hamming weight computation is performed on line 4. To find the violated cut, we perturb the f vector in one coordinate. The coordinate to perturb corresponds to the v_i that is closest to the midpoint of the unit interval. This coordinate is identified on line 5 and f is perturbed accordingly to make it of odd weight on line 6.

Once the possibly violated cut is known, a similarity transform applied on line 8 transforms v to $\tilde{v}$. This aligns the identified cut with the (centered) probability simplex. This transformation is illustrated in Fig. 3 where v is the dot in Fig. 3a, $\tilde{v}$ is the dot in Fig. 3b, and $f = [1\,1\,1]$ (or $f = [0.5\,0.5\,0.5]$). The transformed point $\tilde{v}$ is then projected onto the (centered) probability simplex, as illustrated in Fig. 3c. After projection, the similarity transform is inverted on line 13. The similarity transform is self-inverting.

The execution path up through line 14 of Alg. 2 produces a projection onto the boundary or "shell" of the parity polytope. Through these steps, a point already inside the parity polytope, instead of being left unperturbed, would be projected onto the cut corresponding to the closest odd-weight vertex. To avoid this, we test for parity polytope membership on line 15.

Alg. 3. Given a vector $v \in \mathbb{R}^d$, calculate its Euclidean projection onto the centered probability simplex. $w = \prod_{\mathbb{S}_d - \frac{1}{2}}(v)$ is returned from the method.

1: $\rho \leftarrow \text{descendingSort}(v)$
2: **for** $i \in [d]$ **do** ▷ Calculate Possible Shifts
3: $\quad u_i \leftarrow \frac{1}{i}\left(\sum_{j=1}^{i} \rho_j - 1\right)$
4: **end for**
5: $i^* \leftarrow \max\{i \in [d] : \rho_i > u_i\}$ ▷ Choose Shift
6: **for** $i \in [d]$ **do** ▷ Perform Shift and Clip
7: $\quad w_i \leftarrow \max\left\{v_i - u_{i^*} - \frac{1}{2}, -\frac{1}{2}\right\}$
8: **end for**

We now describe a test for parity polytope membership. If the vector being tested is in the unit hypercube, we need only to check the previously identified cut [32]. Line 15, originally given in [31], tests the hypercube projection of v against the identified cut. This is done by taking the hypercube projection of $\tilde{v}$ and checking on which side of the (centered) probability simplex it lies. If the hypercube projection of v is in the parity polytope, then this must be the parity polytope projection of v since the parity polytope is a subset of the hypercube. A point already in the parity polytope will be left unperturbed.

C. Simplex Projection

We now consider the final important algorithm: Alg. 3, projection onto the centered probability simplex. The centered probability simplex $\mathbb{S}_d - \frac{1}{2}$ is defined by subtracting the all-$1/2$ vector from the probability simplex $\mathbb{S}_d = \left\{v \in \mathbb{R}^d : 1^\top v = 1,\ v_i \geq 0\ \forall\ i \in [d]\right\}$. Projection of $v \in \mathbb{R}^d$ onto the centered probability simplex is a quadratic program. Additionally, projection onto the centered and non-centered probability simplexes are related in the same manner as for the centered and non-centered parity polytope.

Algorithm 3 presents a simplex projection method from [33], modified to project onto the centered simplex. Indeed, computing the projection is a rather straightforward optimization easily solved through analyzing the Karush-Kuhn-Tucker (KKT) conditions [29], [33]. The KKT conditions tell us that the projection is obtained by shifting v along the all-1s vector, clipping components that fall below 0, and ensuring non-clipped components sum to 1. The shift along the all-ones vector results from the fact that the all-1s vector is orthogonal to the simplex, as shown in Fig. 3c. The clipped components are the most negative components of v, therefore the d possible shifts are computed from a sorted version of v. A smart way to identify the best common shift is developed in [34] and used herein. The magnitude of the common shift is computed on lines 2-5 of Alg. 3. The shift and clip are implemented on line 7.

IV. Hardware Architecture

In this section, we build upon well-known hardware architectures for message-passing decoders in the design of a hardware-based ADMM-LP decoder implementation. We modify the arithmetic kernels in the check and variable processing nodes per the operations outlined in Alg. 1.

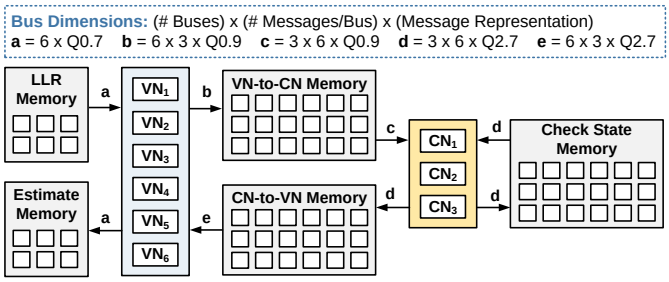

Fig. 4. Partially-parallel architecture for a (3, 6)-regular QC-LDPC code.

A. Decoder

A central challenge in implementing hardware-based decoders is the scalability of the message-passing network. We restrict ourselves to QC codes [35], [36] and a partially-parallel architecture [37]. This simplifies message routing and memory interfacing. We implement the message-passing network with regularly-distributed on-chip FPGA block RAMs. Figure 4 presents an overview of our partially-parallel QC-LDPC decoder architecture for the special case of a $(3, 6)$-regular QC-LPDC code. The architecture is comprised of multiple memory types to store input LLRs, intermediate messages, and output codewords, as well as pipelined Check Node (CN) and Variable Node (VN) processing units that perform the arithmetic operations of Alg. 1.

QC-LDPC codes are defined by a parity-check matrix formed by tilings of $p \times p$ circulant matrices. Each tile of a QC parity-check matrix can be either the all-zeros matrix or some addition of shifted-identity matrices. The tilings naturally divide the parity-check matrix into $s = \frac{n}{p}$ "macro"-columns and $r = \frac{m}{p}$ macro-rows. Inside a given macro-row (column), the required message locations for a check (variable) computation are the locations for the previous check (variable) plus 1 modulo p. This rich class of codes is popular in hardware implementations, appearing in standards such as IEEE 802.11ad [27].

The first step our decoder executes is to load LLRs into memory. We instantiate s memories, each of depth p, to store the LLRs. Each memory is read in parallel to feed LLRs into s pipelined VNs. The VNs also receive messages from a Check-to-Variable (CN-to-VN) message memory, to be discussed later. At the output of the VNs, the current variable estimates x_i are written in parallel into s estimate memories, to be read from upon termination. Variable estimates are also written into Variable-to-Check (VN-to-CN) message memories. There is a VN-to-CN message memory for each shifted-identity matrix in the specification of the parity-check matrix. These memories are addressed using their corresponding shift number to ensure the messages are passed to the proper CN.

Next, r pipelined CNs read their required messages in parallel from the VN-to-CN message memory. In addition, check states are read from check state memories, which are instantiated in the same manner as the VN-to-CN message memories. However, address shifting is not required since these memories are only accessed by CNs. When a CN computation completes, the new check states are written into

check state memories and the messages are written into CN-to-VN message memories. The CN-to-VN message memories are structured in the same manner as the VN-to-CN memories, with write operations using cyclic shift information. The process repeats until the maximum number of iterations is exceeded, or some early termination condition is satisfied.

We find our current implementation of the ADMM-LP decoder to be sensitive to fixed-point quantization. Min-Sum decoders can be implemented with 5- or 6-bit message widths while suffering minimal degradation in error-rate performance compared to floating-point [38]. ADMM-LP requires larger bit-widths. We believe the higher precision is required because the result of the projection operation that CNs perform must be quantized. The quantization results in a loss of precision and a corresponding deterioration of message resolution.

We now discuss the logic that underlies the choices we made in selecting fixed-point representations. We first note that a change in the assignment of bits between integer and fraction parts of fixed-point LLRs amounts to a linear scaling of the LP objective. However, any scaling of the objective in an LP (i.e., of γ in (6)) does not change the solution of the LP. This provides some flexibility in choosing the fixed-point representation of the LLRs. Next, we note that each message passed to a VN can be thought of as either trying to overcome the channel information or as trying to reinforce it. Thus, we allocate any extra bit-width to the integer part of a CN-to-VN message. This provides the dynamic range required to override channel LLRs. In contrast, extra bit-width allocated to VN-to-CN messages should be in the fractional part. An increase in the number of fraction bits mitigates the effect of the inexact (due to finite precision) normalization by $|\mathcal{N}_{v(i)}|$ in the VNs.

Based on this intuition, we select fixed-point message representations to retain as much channel information as possible. We first consider the bit-width of LLRs and the estimate outputs. Respectively, these correspond to the decoder's input and output message widths. Next, we consider how many additional bits VN-to-CN and CN-to-VN messages will receive. VN-to-CN messages, as well as the estimates, lie in the centered unit hypercube. Therefore, these messages receive one sign bit and no integer bits. Next, we give LLRs one sign bit, zero integer bits, and allocate the remainder to fraction bits. This ensures that all channel information is visible in the estimates and the VN-to-CN messages. Experimentation shows that this fixed-point LLR representation provides the best error-rate performance [29]. The CN-to-VN messages are given one sign bit and the same number of fraction bits as the LLRs. This is done so that the summation in the VN computation produces an output that does not have any constant bits for some given LLR. Finally, the check states are given the same representation as the CN-to-VN messages because they are computed in a similar manner.

B. Variable Node

A VN executes lines 4–6 of Alg. 1. An implementation schematic for a VN, labeled with variables from Alg. 1, is depicted in Fig. 5. A VN starts by summing the incoming messages $m_{\mathcal{N}_v(i)\to i}$ and subtracting the LLR γ_i. This large

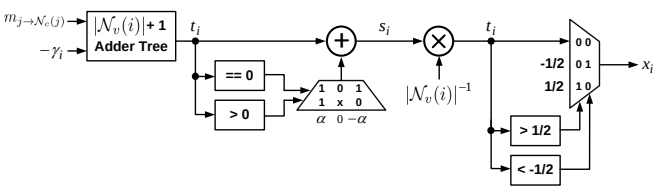

Fig. 5. Variable node compute module.

addition operation is performed using a pipelined adder tree with $\lceil \log_2 (|\mathcal{N}_v(i)| + 1) \rceil$ adder stages with pipelining in between. The output of the adder tree is provided an additional $\lceil \log_2 (|\mathcal{N}_v(i)| + 1) \rceil$ integer bits to prevent overflow.

To implement penalization, the VN checks to see if the adder tree result t_i is greater than, equal to, or less than 0. Using this information, two multiplexers then choose to add α, 0, or $-\alpha$ to the adder tree output. An integer bit is added to the fixed-point representation to avoid overflow.

The next step in the VN is to normalize the penalized sum s_i. Division is generally an expensive operation to perform, but variable degrees are constant for a given code. Therefore, division by $|\mathcal{N}_v(i)|$ can be performed by finding its reciprocal during synthesis and executing the normalization with a multiplication. The fixed-point representation of the reciprocal has 1 sign bit and no integer bits. Our FPGA implementation often uses an on-FPGA Digital Signal Processing (DSP) block to execute this multiplication. Twenty-five bits are used to represent the reciprocal as this is the maximum width accepted by the DSP modules on the FPGA used for our error-rate simulations. Theoretically, this results in a large bit-width for the normalization output, however, unused bits are trimmed during synthesis.

This normalization is trivial for certain variable degrees. For example, if $|\mathcal{N}_v(i)|$ is a power of 2, the normalization can be implemented by bit-shifting the fixed-point representation. Similarly, if the reciprocal of $|\mathcal{N}_v(i)|$ has few ones in its fixed-point representation, soft logic can efficiently implement the resulting multiplication. Thus, to simplify the normalization step, a hardware-oriented code design approach can be taken where $|\mathcal{N}_v(i)|$ is chosen to be a power of 2.

To form the VN output x_i, the above normalization must be projected onto the centered unit interval. Similar to the penalization step, the VN tests whether or not the normalized estimate is less than $-\frac{1}{2}$, greater than $\frac{1}{2}$, or between $-\frac{1}{2}$ and $\frac{1}{2}$. Two multiplexers are used to set the variable estimate to be $-\frac{1}{2}$, $\frac{1}{2}$, or the normalized estimate, respectively.

The final step of the VN architecture is to format the variable estimate x_i to the correct fixed-point representation. The VN-to-CN messages generally have a smaller bit-width than the projected estimate. Since the projected estimate is guaranteed to be between $-\frac{1}{2}$ and $\frac{1}{2}$, its fixed-point representation has 1 sign bit and no integer bits. Therefore, only excess fraction bits need to be removed, which causes the previously mentioned bit trimming for the normalization output. While not indicated in Fig. 5, it is very important to round (rather than truncate) in order to remove these fraction bits. Truncation (i.e., always rounding down) biases decoding towards lower-

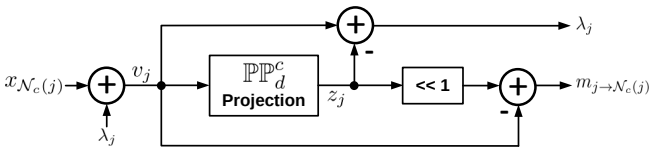

Fig. 6. Check node compute module.

weight codewords. Rounding prevents such a bias.

From this description, one can observe that ADMM-LP VNs are simple to implement. The most complex operation is the adder tree, which gives ADMM-LP VNs $\mathcal{O}\left(|\mathcal{N}_v(i)|\right)$ area scaling and $\mathcal{O}\left(\log|\mathcal{N}_v(i)|\right)$ delay scaling. Additionally, no information needs to be stored in the VN for use in future iterations. The result is a pipeline-friendly module.

C. Check Node

Figure 6 presents a schematic of a CN, which executes the operations on lines 9–12 of Alg. 1. A CN first performs length-$|\mathcal{N}_c(j)|$ vector addition of the incoming message vector $x_{\mathcal{N}_c(j)}$ with the check state vector λ_j. The VN-to-CN messages and check states have the same bit-width, but their fixed-point representations are different. Therefore, to perform the vector addition, check states must be zero extended to have the same number of fraction bits as the incoming messages. The length-$|\mathcal{N}_c(j)|$ vector addition output v_j has components with the same fixed-point representation as the extended check states, except an additional integer bit is added to prevent overflow.

The vector addition output is fed into the parity polytope projection module. It must also be temporarily stored while the projection takes place so it can be used to calculate CN outputs. Implementation of the parity polytope projection, the most resource intensive part of the CN, will be covered in the next subsection. The replica variable vector z_j is assigned the output of the projection module. The replica variable vector has the same bit-width as the projection input, but its fixed-point representation has 1 sign bit and no integer bits since its components are guaranteed to be in $[-\frac{1}{2}, \frac{1}{2}]$.

Following parity polytope projection, new check state values and outgoing messages $m_{j\to\mathcal{N}_c(j)}$ are calculated in parallel using vector addition operations. Before the check state update, extra fraction bits are added to the vector addition result v_j. For the outgoing message calculation, the extended v_j is used with one fewer fraction bit since the parity polytope projection is multiplied by 2. This multiplication is accomplished by bit-shifting the fixed-point representation of z_j.

The final step to output λ_j and $m_{j\to\mathcal{N}_c(j)}$ from the CN is to format their fixed-point representations. To discard excess integer bits, values are saturated at the maximum or minimum that their representations allow. Rounding is performed to discard excess fraction bits. This avoids the aforementioned truncation-induced codeword biases in error-rate performance.

Excluding projection onto the parity polytope, CNs are simple to implement. The vector addition operations have constant delay in check degree and $\mathcal{O}\left(|\mathcal{N}_c(j)|\right)$ area scaling. CN complexity lies in projection onto the parity polytope. Projection gives CNs $\mathcal{O}\left((\log|\mathcal{N}_c(j)|)^2\right)$ delay scaling and $\mathcal{O}\left(|\mathcal{N}_c(j)|\,(\log|\mathcal{N}_c(j)|)^2\right)$ area scaling. Storage of the v_j while the parity polytope projection takes place occupies $\mathcal{O}\left(|\mathcal{N}_c(j)|\,(\log|\mathcal{N}_c(j)|)^2\right)$ area resources with $\mathcal{O}\left((\log|\mathcal{N}_c(j)|)^2\right)$ pipeline stages.

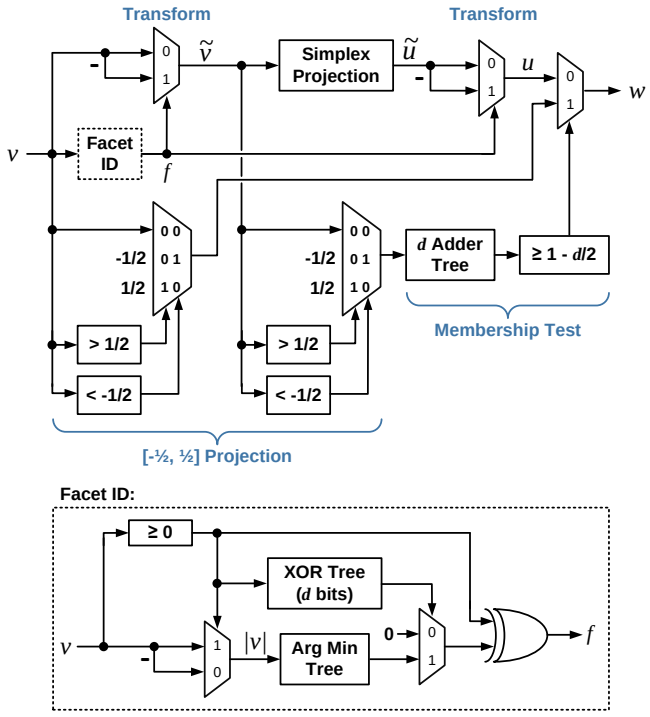

Fig. 7. Parity polytope projection. Note that the (dotted) facet identification module from the upper sub-figure is detailed in the lower sub-figure.

D. Parity Polytope Projection

Figure 7 presents a schematic of the parity polytope projection module. As specified in Alg. 2, the operation starts with facet identification. Finding the closest unit hypercube vertex to the input vector v is accomplished by checking the sign bit of the fixed-point representation of v. Finding the closest odd-weight vertex is slightly more difficult. First, a length-d vector corresponding to the absolute values of v is created with d multiplexers choosing between v_i and $-v_i$ for each component of v. This vector has the same fixed-point representation as v except the sign bit can be dropped since its components are all non-negative. This vector is fed into a min tree that finds the minimum component and outputs a one-hot vector indicating the index of the minimum. If the closest vertex to v is even weight, the one-hot vector is used to flip the bit of f corresponding to the minimum absolute value component via a component-wise XOR. The complexity of this operation lies in the min tree which has $\mathcal{O}\left(\log d\right)$ delay and $\mathcal{O}\left(d\right)$ area scaling. However, v must be stored for $\mathcal{O}\left(\log d\right)$ pipeline stages, resulting in $\mathcal{O}\left(d\log d\right)$ area usage.

With the active facet identified in f, a similarity transform is executed on v to align the active facet with the probability simplex. This is accomplished with d multiplexers choosing between v_i or $-v_i$ based on the value of f_i. The resulting vector $\tilde{v}$ uses the same fixed-point representation as v, however,

since its components are guaranteed to be negative, the sign bit can be dropped and added back in later when required for computation. This operation has constant delay and linear area scaling in the dimensionality of projection d.

At this algorithmic juncture there are three operations that can take place in parallel: projection onto the unit hypercube, projection onto the probability simplex, and testing parity polytope membership. However, our implementation does not execute these operations in parallel. Parallel execution requires knowing the depth of each operation in order to pad properly the lower latency operations with pipeline registers. This can not be done without knowing code check degrees a priori.

Projection of v onto the unit hypercube is the simplest operation to perform. For each component of v, two multiplexers choose between $-\frac{1}{2}$, $\frac{1}{2}$, and v_i. The fixed-point representation is formatted to match the bit-width of the projection output with 1 sign bit and zero integer bits. This operation has constant delay and linear area scaling in d.

Testing parity polytope membership involves projecting $\tilde{v}$ onto the unit hypercube. Hypercube projection is performed in the same manner as above. This is followed by summing the resultant vector using a minimum-depth adder tree. In the adder tree, extra integer bits are added to prevent overflow. By comparing the adder tree result to a constant, we are able to determine what the projection output should be. This decision is stored with a single bit. Due to the adder tree, this operation has $\mathcal{O}(d)$ area scaling and $\mathcal{O}(\log d)$ delay scaling.

The implementation of simplex projection is the topic of the next subsection. Simplex projection dominates the complexity of parity polytope projection. It gives the parity polytope projection $\mathcal{O}\left(d(\log d)^2\right)$ area scaling and $\mathcal{O}\left((\log d)^2\right)$ delay scaling. Additionally, the hypercube projection of v and the active facet identifier f must be stored for the $\mathcal{O}\left((\log d)^2\right)$ pipeline stages it takes to execute the simplex projection. This uses $\mathcal{O}\left(d(\log d)^2\right)$ area. The similarity transform is applied again to the output of the simplex projection to invert itself.

The stored bit indicating parity polytope membership then drives d multiplexers that choose to output the hypercube projection of v or the transformed output of the simplex projection module. Both of these possible outputs have fixed-point representations with 1 sign bit and no integer bits.

E. Simplex Projection

Our algorithm for simplex projection is detailed in Alg. 3; a schematic is depicted in Fig. 8. The components of the vector to be projected are first sorted in descending order. To sort in hardware, we require the set of operations executed to be performed regardless of the input vector. Sorting networks accomplish this. Sorting networks are composed of compare-swap modules, each of which can be implemented with a compare operation and two multiplexers. We implement delay-optimal sorting networks from Knuth [39].

The next step of simplex projection is to calculate all partial sums (termed "prefix sum") of the sorted vector ρ. Since we need to subtract 1 from every partial sum, we simply include -1 as part of the prefix sum input. The prefix sum operation can be performed with $\mathcal{O}(d)$ area scaling and $\mathcal{O}(\log d)$ delay

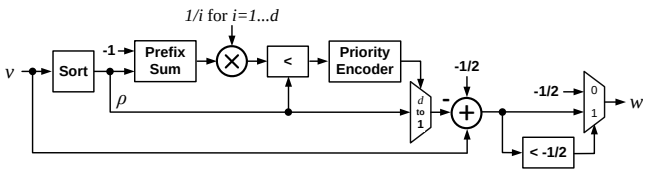

Fig. 8. Simplex projection.

scaling. Ladner and Fischer describe such a construction [40]. The d^{th} sum is computed with a minimum-depth adder tree. Other sums are calculated by reusing computations when possible, making linear area scaling possible. Extra integer bits are allocated to the fixed-point representation of the prefix sum output to prevent overflow. Note that both v and ρ must be stored during this operation, requiring $\mathcal{O}(d \log d)$ area.

Next, the prefix sum output vector components are normalized by their respective component indices. Component index reciprocals are found during synthesis and the normalization is performed by multiplication. As with VNs, multiplication by a power of 2 can easily be implemented with FPGA soft logic for some indices, while a multiplier DSP core is required for others. This operation has constant delay and linear area scaling in the dimension of projection.

We wish to select the normalized partial sum with the largest index that satisfies $\rho_i > u_i$ as the common shift in the simplex projection. First, a length-d binary vector is created indicating the indices satisfying $\rho_i > u_i$. A priority encoder is then used to create a one-hot vector indicating the largest index position satisfying $\rho_i > u_i$. This is also a prefix operation, which yields the same complexity as the prefix sum [21]. However, the operation is on a binary vector, and not a fixed-point vector. The resources consumed are thus much smaller. This one-hot vector is used to select the corresponding component of u.

Finally, the selected component u_{i^*} and $\frac{1}{2}$ are subtracted from all components of v. This is accomplished with two adders for every component of v. After this, each component is compared to $-\frac{1}{2}$ and a multiplexer chooses between the adder results and $-\frac{1}{2}$ to form the final output. The output bit-width matches the input bit-width. The output fixed-point representation has 1 sign bit and no integer bits since each component is between $-\frac{1}{2}$ and $\frac{1}{2}$.

V. Results

In order to test the hardware viability of ADMM-LP decoding, we use an FPGA-in-the-loop simulation environment that consists of a PCI-based Xilinx Virtex-5 FPGA platform on Personal Computer (PC). The proposed architecture was synthesized on the FPGA, along with the wrapper logic needed for noise generation and data transfer to a software test bench.

The binary-input Additive White Gaussian Noise (AWGN) channel is simulated using a Gaussian random number generator [41] on the FPGA. The core is a linear feedback shift register of period 2^{176}, fed into an approximation of the inverse cumulative distribution function. Channel simulation was performed on the FPGA to minimize simulation time by eliminating the bottleneck of PC-to-FPGA data transfer. We verified that FPGA-based channel simulation produced the

$$\begin{bmatrix} 30 & 29 & 27 & 23 & 15 \\ 26 & 21 & 11 & 22 & 13 \\ 6 & 12 & 24 & 17 & 3 \end{bmatrix} \begin{bmatrix} 115 & 13 & 25 & 166 & 17 & 129 \\ 124 & 38 & 137 & 13 & 160 & 136 \\ 75 & 152 & 89 & 73 & 0 & 145 \end{bmatrix}$$

$$\begin{bmatrix} 29 & 30 & 0 & 8 & 33 & 22 & 17 & 4 & 27 & 28 & 20 & 27 & 24 & 23 & - & - \\ 37 & 31 & 18 & 23 & 11 & 21 & 6 & 20 & 32 & 9 & 12 & 29 & 10 & 0 & 13 & - \\ 25 & 22 & 4 & 34 & 31 & 3 & 14 & 15 & 4 & 2 & 14 & 18 & 13 & 13 & 22 & 24 \end{bmatrix}$$

Fig. 9. Shifts for QC parity check matrices where "$-$" denotes an all-zeros matrix. Clockwise from top left: Tanner code, (3, 6) ensemble, WiGig code.

same Frame Error Rate (FER) results as CPU-based channel simulation for low-SNR channels.

Three QC-LDPC codes are considered for error-rate simulation and resource usage analysis. The first is the [155, 64, 20] Tanner code [26], whose parity-check matrix is composed of 31×31 cyclic matrices. The second is the [672, 546] WiGig code [27] composed of 42×42 matrices. The final codes are an ensemble of five [1002, 503] (3,6)-regular QC-LDPC codes. The five parity-check matrices for this ensemble were created by randomly generating shifts for 167×167 identity matrices. The resulting factor graph girths were verified using techniques from [42]. Codes with girth less than 6 were discarded. Example shift matrices are provided in Fig. 9.

Before decoders for these codes are implemented, design decisions regarding the fixed-point representation of messages must be made. The first decision is the input/output bit-width of the decoder. An input/output bit-with of 8 bits was required to guarantee error-rate performance close to double-precision implementations [29]. Fewer bits can be used at the cost of deteriorating error-rate performance. However, the rate of deterioration depends on the code. For example, we found the WiGig code FER performance to be extremely sensitive to decreasing bit-width. Next, the number of bits used for internal messages needs to be decided. Recall that the main effect of these bits is to provide CN-to-VN messages the additional dynamic range needed to override channel information. It was found that 2 additional bits are required to provide good performance in higher-reliability channels [29]. Next, the bit allocations for LLRs, CN-to-VN messages, and check states must be determined. In our architecture, these three values all use the same number of fraction bits. Experimentation indicates that maximizing the number of fraction bits results in the best error-rate performance [29]. That is, LLRs should have no integer bits, and the CN-to-VN messages and check states should have 2 integer bits. We use these allocations in all our simulations.

Table I summarizes the fixed-point precision of each message variable as determined through experimentation to obtain FER performance close to double-precision. Message variables are grouped by computation module and expressed as signed fixed-point numbers in the Q format [43].

An additional parameter that affects error-rates and resource utilization is the number of decoding iterations. Similar to BP, ADMM-LP can be configured to terminate after a maximum number of iterations. This can enforce latency and throughput constraints. Experimentation found that at least 60 iterations are required for our fixed-point configuration to achieve error-rate performance close to its capabilities without a limit on the maximum number of iterations [29].

TABLE I
FIXED-POINT MESSAGE CHARACTERISTICS (MSG / BIT-WIDTH / FORMAT)

Variable Node			**Simplex Projection**		
γ_i	8	Q0.7	v	14	Q4.9
t_i	12	Q4.7	ρ	14	Q4.9
s_i	13	Q5.7	w	14	Q0.13
Check Node			**Parity Polytope Projection**		
x_i, $x_{\mathcal{N}_c(j)}$	10	Q0.9	v	13	Q3.9
$m_{\mathcal{N}_v(i)\to i}$, $m_{j\to\mathcal{N}_c(j)}$, λ_j	10	Q2.7	$\tilde{v}$	14	Q4.9
v_j	13	Q3.9	$\tilde{u}$	14	Q0.13
z_j	13	Q0.12	w	13	Q0.12

A. Error-Rate Performance

The previously mentioned parameter choices affect both error-rate performance and resource consumption. There are two additional parameters that only affect error-rate performance. We discuss the choice of these parameters here.

Simulated channel outputs need to be saturated at some value in order to produce LLRs within the decoder's input range. We parameterize this in terms of standard deviations of channel noise. That is, the channel output is saturated at $\pm(1 + a\sigma)$ where $a > 0$ and σ is the standard deviation of the added Gaussian noise. Experimentation revealed our implementation is not extremely sensitive to this parameter, but $a = 1$ was found to be optimal with respect to FER [29]. Therefore, we saturate channel outputs one standard deviation beyond the transmission values ± 1. The saturated channel outputs are then scaled such that the saturation values are mapped to minimum and maximum LLR values. Recall that AWGN channel outputs are proportional to LLR values, and scaling LLRs does not change the LP decoding objective.

The final parameter configuration is to choose a suitable penalty parameter α. The optimal penalty parameter changes with respect to SNR where larger penalty parameters perform better on low-SNR channels, while smaller penalty parameters perform better on high-SNR channels. A penalty parameter of $\alpha = 0.1$ was found to give good performance across the tested channels for all three codes [29].

Figure 10 presents the FER experimental results for the three codes under investigation on the binary input AWGN channel. We present results for both penalized ($\alpha = 0.1$) and unpenalized ($\alpha = 0$) ADMM-LP , where the value of α refers to its setting in Alg. 4 after the reparameterization that eliminated μ. Double-precision Sum-Product BP and ADMM-LP results are also plotted to form a basis for comparison. Each point on the following plots represents an accumulation of 100 frame errors. Double-precision simulations for ADMM-LP and BP were performed using Liu's implementation [44]. The BP results shown were generated with Butler and Siegel's non-saturating version described in [45]. The same limit of 60 iterations is used for all decoding algorithms.

1) Tanner Code: Fig. 10a presents the FER performance of the Tanner code. A small performance gap exists between the fixed-point and double-precision ADMM-LP implementations. At higher SNRs, all ADMM-LP implementations outperform double-precision BP. The penalized ADMM-LP

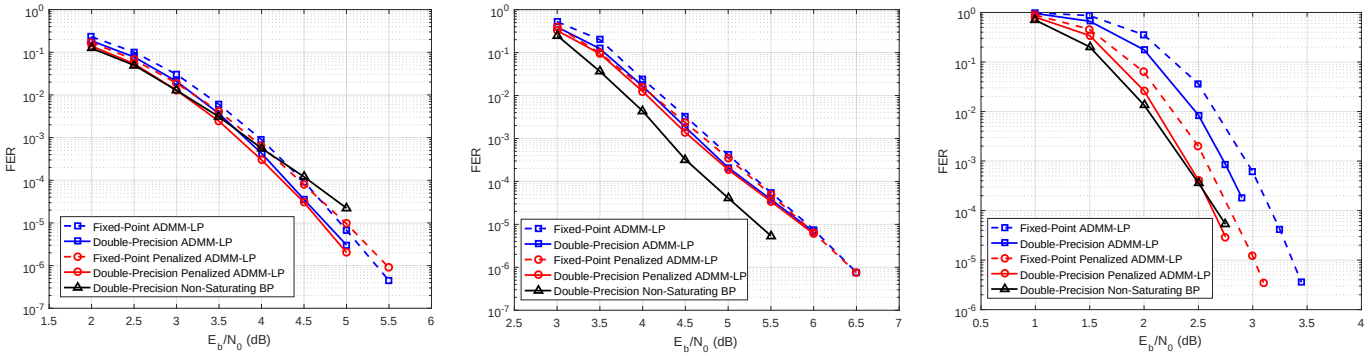

(a) Tanner code (b) WiGig code (c) QC-LDPC ensemble

Fig. 10. FER performance of the Tanner, WiGig, and QC-LDPC ensemble codes.

decoder closes the gap to double-precision BP. However, it does not perform as well as unpenalized ADMM-LP at $E_b/N_0 = 5.5$dB. These results support the conclusion that unpenalized ADMM-LP is better suited to high-SNR channels.

2) WiGig Code: Fig. 10b displays the FER performance of the WiGig code. Again, we see that in both cases, fixed-point ADMM-LP maintains very close performance to the double-precision implementation with 10-bit messages. However, there is a very large performance gap between BP and ADMM-LP. This performance gap is not closed with the addition of penalization; the opposite of what has been observed with other codes [15]. The root cause of the weakness of LP decoding for this code requires further investigation. We conjecture that, in part, it may be due to the high-degrees of the check nodes, resulting in more pseudocodewords, or due to the variable-one variable nodes.

3) QC-LDPC Ensemble: The FER performance of the ensemble of $(3, 6)$-regular QC-LDPC codes is shown in Fig. 10c. Each curve is obtained by averaging the performance of the same five codes from the QC-LDPC ensemble. This experiment is a more powerful demonstration of the performance of ADMM-LP, where the addition of penalization closes the large performance gap between BP and ADMM-LP. The fixed-point implementations of both penalized and unpenalized ADMM-LP achieve performance very close to double-precision.

B. Resource Usage

This section examines the FPGA resource utilization for the three decoders synthesized for the Tanner, WiGig, and QC-LDPC ensemble codes using the fixed-point message representations summarized in Table I. While the error-rate simulations were performed on an older Xilinx Virtex-5 FPGA, the resource utilization results presented here target a newer Altera Stratix V FPGA (model 5SGXEA7N2F45C2). This FPGA has 234,720 Adaptive Logic Modules (ALM), 256 DSP blocks, and 2,560 "M20K" Random Access Memory (RAM) blocks with 52,428,800 RAM bits in total. Synthesis was performed using Altera's Quartus II (15.0.0) tool suite using balanced optimization. Basic power estimation was performed using gate-level simulations of high-noise decodings and Altera's power analyzer tool. Table II summarizes the FPGA utilization and throughput results.

TABLE II
FPGA PERFORMANCE RESULTS AND RESOURCE COMPARISON

Specification	Tanner	WiGig	Ensemble
Block Length	155	672	1002
Code Rate	2/5	13/16	1/2
QC Matrix Structure (r rows, s columns)	$r = 3$ $s = 5$	$r = 3$ $s = 16$	$r = 3$ $s = 6$
Expansion Factor	31	42	167
FPGA Performance Results			
Clock Freq. (MHz)	237	221	225
Iterations	60	60	60
Cycles Per Iter.	137	189	440
Throughput (Mb/s)	4.47	13.16	8.52
Latency Per Iter. (μs)	0.578	0.851	1.960
Total Est. Power (mW)	758	2153	863
Dynamic Power (mW)	703	1942	797
Static Power (mW)	55	211	66
Adaptive Logic Modules	12201	34715	14315
DSP Blocks	11	46	15
Memory (Bits)	16,430 (40 BRAMs)	67,284 (92 BRAMs)	106,212 (47 BRAMs)
FPGA ALM Resource Utilization			
Check Nodes	85%	94%	85%
Variable Nodes	5%	3%	6%
Memory	8%	2%	8%
Control Logic	2%	1%	1%
Maximum Number of Pipeline Stages			
Check Node	46	54	47
Variable Node	9	9	9
FPGA Power Breakdown			
Check Nodes	76%	84%	76%
Variable Nodes	7%	4%	6%
Memory	17%	12%	17%

1) Tanner Code: The implementation of the Tanner code decoder in the partially-parallel architecture has three degree-5 CNs and five degree-3 VNs. Each VN uses a single DSP block to accomplish its normalization, and each CN uses two DSP blocks to perform division by 3 and division by 5.

From Table II, we see that CNs account for the majority of resource usage. Therefore, a further breakdown of CN resource consumption is warranted. We now break down the ALM and power consumption inside a CN on a sub-component basis. Parity polytope projection accounts for 89% of ALM and 91% of power usage inside the check node. Simplex projection accounts for a bit over 51% and 54%, respectively. Finally, sorting consumes 14% of ALM and power usage, and prefix addition consumes 11% of ALM and power usage. Note that

TABLE III
COMPARISON OF FPGA-BASED LDPC DECODER IMPLEMENTATIONS

	[46] 2006	[47] 2007	[48] 2005	[49] 2011	This Paper
Alg	BP	Min-Sum	N/A	Min-Sum	ADMM-LP
Arch	Serial	FP	PP	PP	PP
Length	980	1200	1038	1152	1002
Dsn Rate	0.696	1/2	1/2	1/2	1/2
Struct	N/A	PEG	QC	QC	QC
Max Iter	100	10	18	10	60
Perf @ 3dB	BER 8×10^{-5}	FER 5×10^{-2}	N/A	BER 1×10^{-5}	FER 1×10^{-5} BER 2×10^{-7}
Device	AC-EP1C6	XV-4	XV-E	XV-2	AS-V
Msg Bit Width	6	3	4	4	LLR 8 Int. 10
Early Term	No	No	No	No	No
Freq (MHz)	136	100	26	64	224
Thpt (Mb/s)	7	6000	72	50	8.52
Thpt/Iter (Mb/s)	0.07	600	4	5	0.142
Delay / Iter (μs)	1.40	0.02	0.80	2.30	1.96
Pwr (mW)	N/A	N/A	322	N/A	863
Resources	997 ALMs	40613 Slices	10883 Slices	2778 Slices	14315 ALMs
Mem (Kbits)	34	N/A	N/A: 120 BRAMs	19.5Kb: 29 BRAMs	106.2Kb: 47 BRAMs

Note that the BER presented here corresponds to that achieved by fixed point penalized ADMM-LP at FER $= 1.2\times10^{-5}$ as plotted in Fig. 10c. The acronyms used are as follows: FP = "fully parallel", PP = "partially parallel", PEG = "progressive edge growth", AE = "Altera Cyclone", XV = "Xilinx Virtex", AS = "Altera Stratix".

these figures are nested. For example, the 91% of CN power usage attributed to parity polytope projection includes power used in simplex projection. We believe this is due to heavy ALM usage for intermediate storage. For example, inside a CN, v_j must be stored until projection is complete. In our implementation, the resources used for this storage count toward polytope projection resource consumption. Since area utilization and power are related, it is not surprising that the ALM and power breakdowns are quite similar.

2) WiGig Code: The WiGig code has one degree-16 CN, one degree-15 CN, and one degree-14 CN. It has 14 degree-3 VNs, one degree-2 VN, and one degree-1 VN. The degree-3 VNs use fewer resources than the Tanner decoder due to increased resource sharing among the 14 degree-3 VNs.

Again, CNs account for the majority of resource usage. The percentage of ALM usage and power consumption inside the degree-16 CN on a sub-component basis are very similar to the degree-5 CNs of the Tanner decoder. For the WiGig code, the two complexity-dominating operations, sort and prefix addition, consume a larger fraction of resources.

3) QC-LDPC Ensemble: The QC-LDPC ensemble implementation has six degree-3 VNs and three degree-6 CNs. CNs account for the majority of resource usage, and the internal CN resource breakdown is again almost identical to that of the Tanner code CNs. The same trend has emerged for all decoders, where pipeline depth and intermediate value storage have a large impact on resource consumption.

C. Implementation Comparison

Table III compares our implementation for the QC-LDPC ensemble with several FPGA-based LDPC decoders having similar code rates and comparable block length.

Our ADMM-LP decoder achieves better error-correction performance at the chosen $E_b/N_0 = 3$dB operating point compared to the four comparison works as ADMM-LP outperforms Min-Sum in this metric. On the other hand, our decoder requires more iterations and larger bit widths, resulting in in lower throughput and higher logic utilization.

A direct comparison of logic resources and overall area between designs implemented on Altera and Xilinx FGPAs is nearly impossible. The internal lookup table and D-flipflop structure of a Xilinx Slice is not equivalent to an Altera ALM [50], and the ALM and Slice architectures change from one FPGA generation to another. Additionally, FPGA synthesis and place-and-route stages are highly dependent on the target device. Nevertheless, the logic utilization numbers of Table III show that our partially-parallel ADMM-LP decoder implementation has a logic resource utilization within an order of magnitude of the partially-parallel comparison works implemented on Xilinx devices. As expected, our level of resource utilization is between the comparison implementations whose decoders realize serial and fully-parallel architectures.

VI. DISCUSSION AND CONCLUSIONS

In this paper we demonstrate that ADMM-LP decoding can attain excellent error-rate performance in a fixed-point implementation. While our initial implementation requires higher fixed-point precision and more logic resources than the Min-Sum algorithm, this study points to numerous possible avenues for future developments, which could bring ADMM-LP's resource requirements into line with those of other message-passing decoders.

One avenue is algorithmic simplification. Just as Min-Sum can be viewed as a computationally simple approximation of Sum-Product BP, we can seek approximations of ADMM-LP that preserve its the high-SNR performance. As one example, in [29] it is observed that implementing partial-sort (rather than full-sort) can result in a negligible increase in error rates.

A second set of directions is hardware-centric. Numerous interesting challenges remain in the design of a hardware-efficient implementation. For example, it is not obvious how to implement a CN or a VN unit that can handle multiple node degrees. We believe that this problem can be solved through innovative hardware sharing or algorithmic generalization. As a second example, ADMM-LP also provides an opportunity for simplifying message-passing networks, especially when considering a fully-parallel architecture. This is because the same message is sent from each variable to all connected checks. Such message broadcasting can perhaps be exploited to reduce interconnect complexity. Finally, this study is a first step en-route to the development of a fully custom, in-silicon, Application Specific Integrated Circuit (ASIC). An ASIC would allow for high-performance, power-optimized register files and customized message-passing resources that would yield significant performance improvements not possible in an FPGA realization.

Referring to Table III, we note that while our normalized throughput per iteration is $35\times$ lower than that of the Min-Sum decoder of [49], our ADMM-LP decoder achieves a Bit Error Rate (BER) nearly $100\times$ better. This is the crux of the matter. If one is concerned with applications where excellent performance in the high-SNR regime is required, a regime where algorithms such as Min-Sum or Sum-Product encounter error-floor problems, then ADMM-LP should be an algorithm of great interest. Our current implementation outperforms Min-Sum with less than an order of magnitude difference in the number of FPGA resources required. Further development, and innovation, could turn ADMM-LP into the algorithm of choice in such regimes of operation.

REFERENCES

[1] C. Berrou, A. Glavieux, and P. Thitimajshima, "Near Shannon limit error-correcting coding and decoding: Turbo-codes," in *Proc. IEEE Int. Conf. Comm.*, May 1993, pp. 1064–1070.

[2] D. J. C. MacKay and R. M. Neal, "Good codes based on very sparse matrices," in *IMA Int. Conf. on Cryptography and Coding*, C. Boyd, Ed. Springer, Dec. 1995, pp. 100–111.

[3] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 498–519, Feb. 2001.

[4] T. Richardson, "Error floors of LDPC codes," in *Proc. Allerton Conf. Comm. Control and Comp.*, vol. 41, no. 3, Oct. 2003, pp. 1426–1435.

[5] J. Feldman, M. J. Wainwright, and D. R. Karger, "Using linear programming to decode binary linear codes," *IEEE Trans. Inf. Theory*, vol. 51, no. 3, pp. 954–972, Mar. 2005.

[6] M. H. Taghavi and P. H. Siegel, "Adaptive methods for linear programming decoding," *IEEE Trans. Inf. Theory*, vol. 54, no. 12, pp. 5396–5410, Nov. 2008.

[7] J. Feldman, T. Malkin, R. A. Servedio, C. Stein, and M. J. Wainwright, "Message-passing algorithms and improved LP decoding," in *Proc. Int. Symp. Inf. Theory*, Chicago, IL, Jun. 2005.

[8] A. Arora, D. Steuer, and C. Daskalakis, "Message-passing algorithms and improved LP decoding," in *ACM Symposium on Theory of Computing (STOC)*, May 2009.

[9] P. O. Vontobel and R. Koetter, "Towards low-complexity linear-programming decoding," in *Proc. 4th Int. Symp. Turbo Codes and Related Topics*, Munich, Germany, Apr. 2006, pp. 1–9.

[10] D. Burshtein, "Iterative approximate linear programming decoding of LDPC codes with linear complexity," *IEEE Trans. Inf. Theory*, vol. 55, no. 11, pp. 4835–4859, Nov. 2009.

[11] S. Barman, X. Liu, S. C. Draper, and B. Recht, "Decomposition methods for large scale LP decoding," *IEEE Trans. Inf. Theory*, vol. 59, no. 12, pp. 7870–7886, Dec. 2013.

[12] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Machine Learning*, no. 1, pp. 1–122, 2011.

[13] X. Liu and S. C. Draper, "Instanton search algorithm for the ADMM penalized decoder," in *Proc. Int. Symp. Inf. Theory*, Honolulu, Jun. 2014.

[14] ——, "ADMM decoding on trapping sets," in *Proc. Int. Symp. Inform. Theory*, Hong Kong, Jun. 2015.

[15] ——, "The ADMM penalized decoder for LDPC codes," *IEEE Trans. Inf. Theory*, vol. 62, no. 6, pp. 2966–2984, Jun. 2016.

[16] Y. Wang, J. S. Yedidia, and S. C. Draper, "Multi-stage decoding of LDPC codes," in *Proc. Int. Symp. Inf. Theory*, South Korea, Jul. 2009.

[17] X. Liu and S. C. Draper, "ADMM LP decoding of non-binary ldpc codes in $\mathbb{F}_{2^m}$," *IEEE Trans. Inf. Theory*, vol. 62, pp. 2985–3010, Jun. 2016.

[18] ——, "LP-decodable multipermutation codes," *IEEE Trans. Inf. Theory*, vol. 62, pp. 1631–1648, Apr. 2016.

[19] X. Zhang and P. H. Siegel, "Efficient iterative LP decoding of LDPC codes with alternating direction method of multipliers," in *Proc. Int. Symp. Inf. Theory*, Istanbul, Turkey, Jul. 2013, pp. 1501–1505.

[20] G. Zhang, R. Heusdens, and W. B. Kleijn, "Large scale LP decoding with low complexity," *IEEE Comm. Lett.*, pp. 2152–2155, Nov. 2013.

[21] M. Wasson and S. C. Draper, "Hardware based projection onto the parity polytope and probability simplex," in *Proc. 49th Asilomar Conf. Signals, Systems, Computers*, Pacific Grove, CA, Nov. 2015, pp. 1015–1020.

[22] I. Debbabi, B. L. Gal, N. Khouja, F. Tlili, and C. Jego, "Fast converging ADMM-penalized algorithm for LDPC decoding," *IEEE Commun. Lett.*, vol. 20, no. 4, pp. 648–651, Apr. 2016.

[23] ——, "Analysis of ADMM-LP algorithm for LDPC decoding, a first step to hardware implementation," in *IEEE Int. Conf. Electronics, Circuits, and Systems*, Cairo, Egypt, Dec. 2015, pp. 356–359.

[24] X. Jiao, H. Wei, J. Mu, and C. Chen, "Improved ADMM Penalized decoder for irregular low-density parity-check codes," *IEEE Commun. Lett.*, vol. 19, no. 6, pp. 913–916, Jun. 2015.

[25] H. Wei, X. Jiao, and J. Mu, "Reduced-complexity linear programming decoding based on ADMM for LDPC codes," *IEEE Commun. Lett.*, vol. 19, no. 6, pp. 909–912, Jun. 2015.

[26] R. M. Tanner, D. Sridhara, and T. Fuja, "A class of group-structured LDPC codes," in *Proc. ICSTA*, Ambleside, U.K., Jul. 2001.

[27] *IEEE Std 802.11ad-2012 (Amend. to IEEE Std 802.11-2012, as amended by IEEE Std 802.11ae-2012 and 802.11aa-2012)*, pp. 1–628, Dec 2012.

[28] J. Feldman, "Decoding error-correcting codes via linear programming," Ph.D. dissertation, Massachusetts Institute of Technology, USA, 2003.

[29] M. Wasson, "Hardware-based linear program decoding with the alternating direction method of multipliers," Master's thesis, University of Toronto, Canada, Nov. 2016.

[30] S. Barman, X. Liu, S. C. Draper, and B. Recht, "Decomposition methods for large scale LP decoding," in *Proc. Allerton Conf. Comm. Control Computing*, Monticello, IL, Sep. 2011.

[31] X. Zhang, "LDPC codes: Structural analysis and decoding techniques," Ph.D. dissertation, University of California, San Diego, USA, 2012.

[32] X. Zhang and P. H. Siegel, "Adaptive cut generation algorithm for improved linear programming decoding of binary linear codes," *IEEE Trans. Inf. Theory*, vol. 58, no. 10, pp. 6581–6594, Oct. 2012.

[33] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra, "Efficient projections onto the ℓ_1-ball for learning in high dimensions," in *Proc. Int. Conf. Machine Learning*, San Diego, USA, Dec. 2008.

[34] S. Shalev-Shwartz and Y. Singer, "Efficient learning of label ranking by soft projections onto polyhedra," *Journal of Machine Learning Research*, vol. 7, pp. 1567–1599, Jul. 2006.

[35] Y. Kou, S. Lin, and M. P. C. Fossorier, "Low-density parity-check codes based on finite geometries: A rediscovery and new results," *IEEE Trans. Inf. Theory*, vol. 47, no. 7, pp. 2711–2736, Nov. 2001.

[36] M. P. C. Fossorier, "Quasicyclic low-density parity-check codes from circulant permutation matrices," *IEEE Trans. Inf. Theory*, vol. 50, no. 8, pp. 1788–1793, Aug. 2004.

[37] D. E. Hocevar, "A reduced complexity decoder architecture via layered decoding of LDPC codes," in *Proc. Work. Signal Proc. Sys.*, Oct. 2004.

[38] Y. S. Park, D. Blaauw, D. Sylvester, and Z. Zhang, "Low-power high-throughput LDPC decoder using non-refresh embedded DRAM," *IEEE J. Solid-State Circuits*, vol. 49, no. 3, pp. 783–794, Mar. 2014.

[39] D. E. Knuth, *The Art of Computer Programming: Sorting and Searching*, 2nd ed. Redwood City, USA: Addison Wesley Longman, 1998, vol. 2.

[40] R. E. Ladner and M. J. Fischer, "Parallel prefix computation," *J. of the ACM*, vol. 27, no. 4, pp. 831–838, Oct. 1980.

[41] G. Liu. (2015) Gaussian noise generator. OpenCores. [Online]. Available: http://opencores.org/project,gng

[42] Y. Wang, S. C. Draper, and J. S. Yedidia, "Hierarchical and high-girth QC LDPC codes," *IEEE Trans. Inf. Theory*, vol. 59, no. 7, pp. 4553–4583, Jul. 2013.

[43] S. A. Khan, *Digital design of signal processing systems: a practical approach*. John Wiley & Sons, 2011.

[44] X. Liu. (2015) ADMM decoder. [Online]. Available: https://sites.google.com/site/xishuoliu/codes

[45] B. K. Butler and P. H. Siegel, "Error floor approximation for LDPC codes in the AWGN channel," *IEEE Trans. Inf. Theory*, vol. 60, no. 12, pp. 7416–7441, Dec. 2014.

[46] X. Lei, T. Zhenhui, and Y. Dongping, "The moderate-throughput and memory-efficient LDPC decoder," in *2006 8th International Conference on Signal Processing*, vol. 3, 2006.

[47] R. Zarubica, S. G. Wilson, and E. Hall, "Multi-Gbps FPGA-based low density parity check (LDPC) decoder design," in *in Proc. Global Telecomm. Conf.*, Nov 2007, pp. 548–552.

[48] P. Bhagawat, M. Uppal, and G. Choi, "FPGA based implementation of decoder for array low-density parity-check codes," in *in Proc. Int. Conf. Acoustics, Speech, and Signal Proc.*, Mar. 2005.

[49] V. A. Chandrasetty and S. M. Aziz, "A multi-level hierarchical quasi-cyclic matrix for implementation of flexible partially-parallel LDPC decoders," in *Int. Conf. Multimedia, Expo*, July 2011, pp. 1–7.

[50] "Stratix II vs. Virtex-4 density comparison," Altera White Paper, Aug. 2002. [Online]. Available: https://www.altera.com/en_US/pdfs/literature/wp/wpstxiixlnx.pdf